Executive UPDATE
CUTTER CONSORTIUM
Access to the Experts
BUSINESS & ENTERPRISE ARCHITECTURE

Sense-Analyze-Respond: Lead the Disruption

by Jagdish Bhandarkar, Lakshmanan G, Raghavan S, and Shyam Doddavula

Successful organizations have a dilemma. Should they continue growing what is making them successful, or should they disrupt the market that they are successfully leading? What if they choose to grow their current business and later get disrupted by an upstart innovation? Or what if they disrupt a successful business too early? What if they overestimate the potential of an external disruption?

These are very tough questions that successful organizations face every day, especially in these times as digital transformation has become front of mind. To tackle such dilemmas, we believe that an organization, in addition to being led by able individuals, must provide those individuals with a reasoning framework that poses these questions, answers them, and helps realize the answers with regard to organizational assets.

In this *Executive Update*, we propose a real-time enterprise (RTE) framework for one such sense-analyze-respond framework. Armed with a practical and well-oiled RTE framework, enterprises can minimize their chances of being blindsided by adverse developments, internal or external, and maximize their chances of leading disruption in their markets.

The Real-Time Enterprise Framework

As shown in Figure 1, enterprises need to have a three-step framework consisting of:

1. A **sense** component that enables the enterprise to be informed of all relevant external and internal information

2. An **analyze** component that helps the enterprise analyze the information, identify patterns, and predict likely impacts and consequences

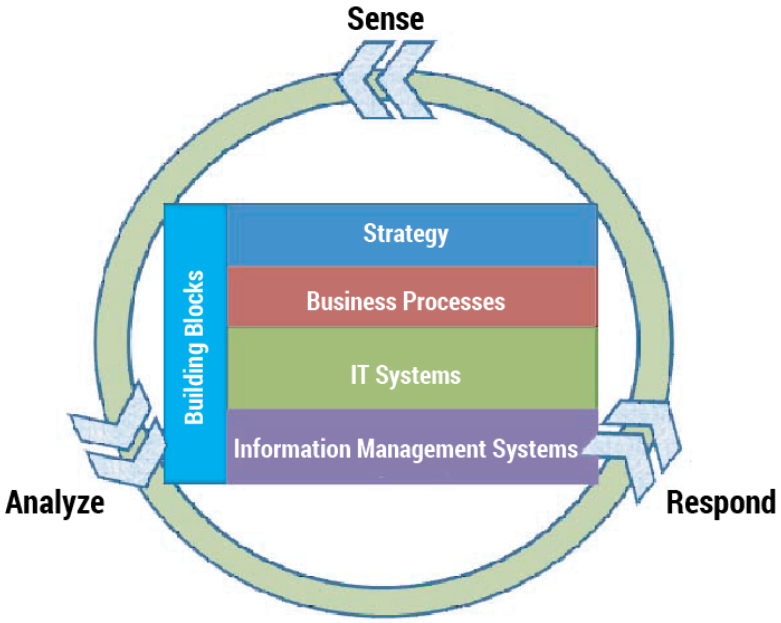

Figure 1 — The RTE framework.

3. A **respond** component that determines the best possible response based on analysis of the relevant information and executes that response

The business strategy of the enterprise determines what needs to be sensed, what kind of analysis is needed, and what the response needs to be, while the enterprise's business processes, information, and IT systems support the collection and analysis of information and the execution of the response.

Sense

The sense component needs to capture the relevant internal and external information and any changes in these conditions that are relevant for the enterprise. External factors include:

- Market conditions and changes in market dynamics
- The competitive landscape
- New threats and opportunities
- Changes such as acquisitions or new regulatory compliance requirements
- New laws that may affect the operating environment
- Customer experiences and brand perceptions

- Partner and customer feedback
- Information from social media

Internal factors include:

- Performance of various products and services
- Health of internal systems
- Insights about system usage
- Feedback and ideas based on collective wisdom of employees

Analyze

The analyze component helps identify the patterns, trends, and correlations across the various external and internal inputs. It enables the enterprise to forecast numerous potential scenarios, identifying the probabilities and consequences with what-if analyses.

Respond

The respond component involves identifying response strategies, which can range from ignoring certain sensed and analyzed events as irrelevant, to making incremental changes in direction with reallocation of resources, to making radical changes such as exiting or entering new businesses or geographies. IT has a key role to play in allowing more response options and enabling fast execution of the chosen responses. The following sections describe how to build technology solutions that enable a real-time enterprise.

Getting Business Processes and IT Systems Ready for RTE

To be able to react to changing business and market conditions quickly, particularly among today's various digital transformations that surface almost daily, enterprises must be agile at various levels. The enterprise needs to decide on the business strategy fast and execute that strategy fast. To do this, both business processes and IT systems must be able to meet the time-to-market requirements or scalability needs of the enterprise.

Enterprises must architect their business and IT systems with the following six characteristics:

1. Agility
2. Anywhere, anytime availability
3. Scalability
4. Intelligence
5. Collaboration
6. Low latency

Agility

To meet time-to-market requirements, business processes should be automated as much as possible, and IT systems should be architected so that the various layers (infrastructure, platform, application) can be provisioned on demand in an elastic fashion to meet the dynamic workloads. Virtualization and service orientation are two key technology strategies to adopt. Moreover, many of today's innovative enterprises are moving toward robotic process automation as a strategy.

Anywhere, Anytime Availability

Business processes need to enable customers, partners, and employees to use enterprise IT systems and be productive from anywhere, anytime. The IT systems should allow secure access over multiple channels and devices, across the global stage.

Scalability

With increasing adoption of social media (and more and more platforms popping up) and everyday devices such as your refrigerator getting instrumented and connected to the Internet, we are experiencing a deluge of data. Next-generation enterprise IT systems need to be able to deal with such large volumes and wide varieties of data. IT systems also need to leverage cloud computing technologies in order to be scalable and elastic enough to store, process, and analyze billions of interactions with connected smart devices.

Intelligence

IT systems should be designed to be continuously learning so that they adapt to changing conditions and usage patterns. They should be able to leverage predictive analytics, artificial intelligence, and machine learning capabilities.

Collaboration

An enterprise's business processes should be designed to leverage the collective wisdom of all its stakeholders, and the IT systems should be designed for collaboration.

Low Latency

Change is rapid, and there is increasing need for real-time processing. Architectural patterns like event-driven architectures and technologies like complex event processing (CEP) enable near-real-time processing and responses. IT systems should be designed to leverage such strategies and technologies.

Getting Information Management Systems Ready for RTE

Data present in various IT systems is an important business asset for organizations. Data-driven enterprises implement processes and systems that make optimal use of this data by making it available to various business processes in a timely and reliable manner. This enables multiple business processes and provides a big competitive advantage to such organizations. Being data-driven is an important enabler for real-time enterprises.

Characteristics

Key characteristics of a data-driven, real-time enterprise are:

- **Timely access to data.** Data should be available in real time or near real time as required by the business processes.
- **Accessible data.** Data should be accessible to all business processes and people through varied channels.
- **Accurate data.** Data should be accurate and reliable.
- **Relevant data.** Data should be relevant to the operational processes and available in the required form (raw data or insights).
- **Bidirectional data flow.** Information should flow in both directions between operational and analytics systems.

Need for Being Data-Driven

There are several business scenarios that require data from multiple sources and at latencies ranging from subseconds to a few seconds, minutes, or hours, depending upon the business need. Examples include:

- Real-time product recommendations and website personalization for customers based on history
- Fraud detection for financial transactions through matching with prior existing patterns
- Promotions optimization based on real-time sales and promotion performance information
- Location-based offers for consumers
- Product pricing optimization based on real-time sales information
- Supply chain optimization based on real-time feedback
- Customer service optimization through creation of a near-real-time profile of the customer
- Equipment failure detection

Enabling these scenarios results in multiple business benefits, such as more satisfied customers, higher sales, lower costs, and faster product launches.

Challenges

There are several challenges that enterprises may have to overcome before they can make the best use of their data and become truly data-driven. Following are some examples:

- Diverse data sources, including multiple systems, multiple data-generating devices, multiple technologies
- Volume of data; that is, a vast number of data elements being generated at a very fast pace
- Variety of data, including structured (transaction data) and unstructured data (social media, emails, documents, etc.)
- Data extraction in real time or near real time without impacting the source systems
- Integrating, filtering, and quality-checking data in real time
- Fast load and query at the same time in analytical systems
- Complex data analysis, including aggregating, mining, forecasting
- High availability of analytical systems due to real-time need of insights
- Delivery through multiple channels (email, messaging, desktop, tablets, smartphones, etc.)

Key Enablers

Following are the key enablers for a data-driven, real-time enterprise:

- **Real-time data integration.** This allows the enterprise to extract data from various operational systems, clean it, integrate it, filter it, and make it available to various operational processes or analytical systems in real time or near real time, as required by the business.
- **Real-time data warehouses (DWs).** Real-time DWs allow analysis of the operational data at very low latencies, ranging from subseconds to a few seconds, minutes, or hours. Real-time DWs are one of the most important enablers for the data-driven enterprise.

Architecture Pattern Using the Key Enablers

Figure 2 represents the architecture pattern for real-time data integration and data warehousing, with its key components and the underlying technologies.

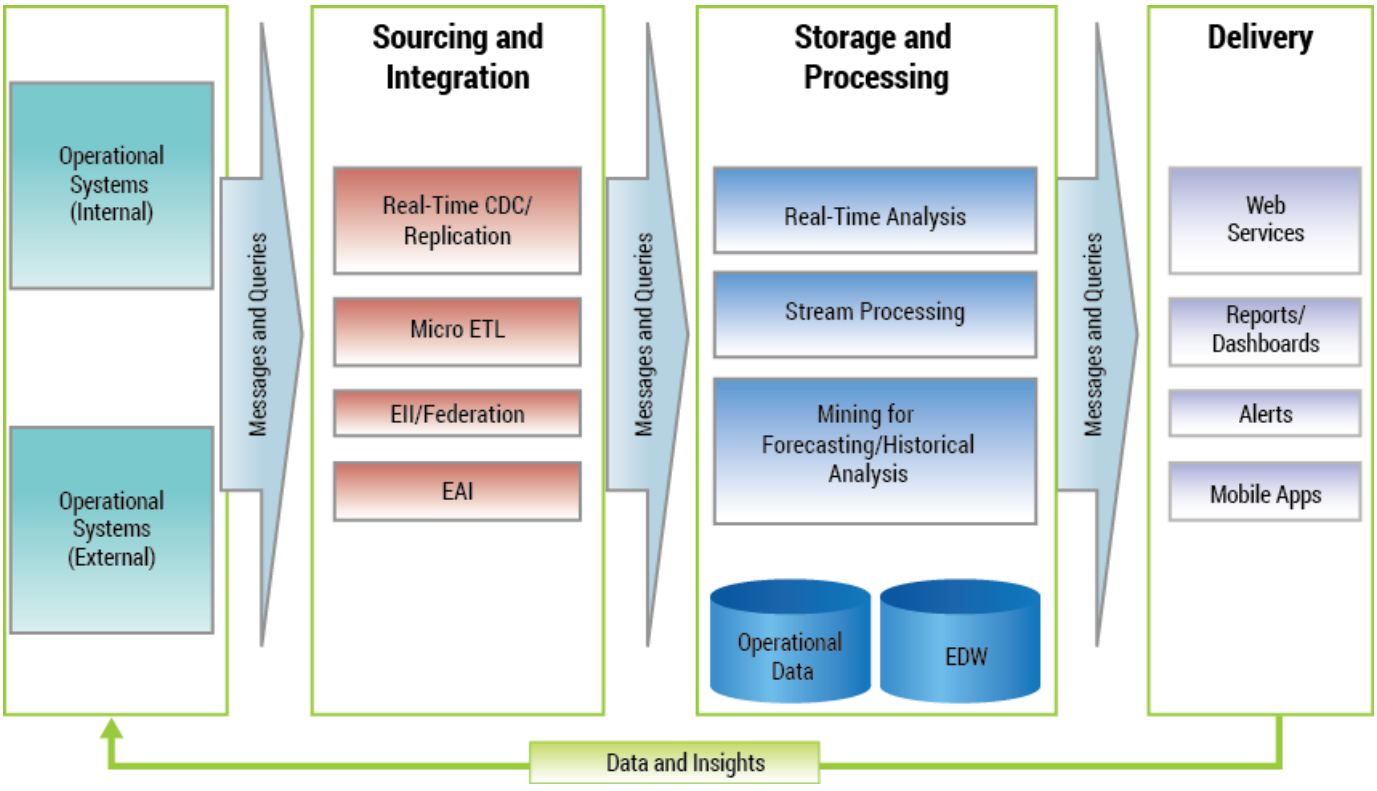

Figure 2 — Architecture pattern for real-time data integration and data warehousing.

Data Delivery

The following techniques help deliver real-time insights to various processes and people:

- **SOA-based analytics.** Access to the analytics is made through Web services.
- **Real-time reports/dashboards.** On-demand or very frequent refresh of the reports/dashboards can enable executives to take decisions based on real-time data.
- **Alerts (emails, texts, messages).** Any anomaly detected during analysis should be sent across to the relevant system or person that can take immediate action to resolve the issue. Real-time analytical systems should interface with the messaging systems, which can deliver alerts through emails, texts, or messages.
- **Mobile apps.** It is imperative for enterprises to enable real-time data analysis through mobile devices. This can be achieved through native apps or HTML5-based apps.

Data Storage and Processing

DW appliances or analytic DBMSs are a class of database management systems designed to handle complex analytical queries on large amounts of data while providing faster performance for queries and data load. These database systems, which mostly support structured data and SQL-based data access, include:

- **Enterprise data warehouse (EDW).** An EDW contains the historical data used for data mining, forecasting, and trend analysis. In some scenarios, if the EDW can handle mixed workloads of real-time and historical data, it can also store the real-time operational data.
- **Operational data store (ODS).** An ODS stores the integrated real-time operational data loaded through one of the data integration methods. An ODS only stores the current data, augmenting the historical data stored in the EDW.
- **Hadoop/HDFS and NoSQL databases.** These databases are suitable for processing unstructured data and doing complex analytical processing. They can be used for building very scalable and cost-effective solutions. Data is stored on a distributed file system (HDFS) or in NoSQL databases. NoSQL databases can be categorized as document stores, key value stores, graph databases, so on.
- **Data mining, forecasting, and historical analysis.** This component provides the traditional business analytics capabilities the business requires. These analyses are made available as on-demand services or precalculated values so that they can be accessed in real time to provide historical context to real-time data.

- **Stream analysis, complex event processing.** CEP is the continuous and incremental processing of event streams from multiple systems for zero-latency response times (as in fraud detection, fault detection, etc.). Real-time events are compared against the historical data patterns pulled from the EDW. Data processing performance is optimized through techniques such as in-memory caching and aggregation over time windows.

- **Real-time/on-demand data processing.** Data analysis is made available through a set of services that can be invoked by consumers depending upon the operational need. These services pull the real-time data through data federation (see below) or from an ODS. They combine the real-time analysis with the historical analysis and patterns, pulled from the EDW through data analysis and mining services.

Data Extraction

There are several technologies that enable real-time data capture from the source systems. These differ in terms of the speed of data capture and their impact on the source systems:

- **Change data capture (CDC).** CDC represents the set of patterns that identify the changes that have occurred over a period of time. This allows the processing of only the changed data, which is usually smaller than the entire data set, and thus helps in processing the data in real time. There are several techniques that enable CDC, including timestamps on source data, database triggers, and database log scrapping. CDC solutions are proven solutions that ensure better performance for the subsequent ETL processes. Data can also be loaded to the operational data warehouses directly.

- **Real-time data replication.** Replication provides data extraction and publishing in real time with minimal latency and minimal impact to the source operational systems. There are several categories of replication tools that support heterogeneous data sources, big data–specific requirements, multi-directional data movement, and API-based replication. Real-time data replication also supports a variety of data sources and targets, such as files and FTP locations. Replication is a proven technology and one of the critical components of real-time technologies.

- **Enterprise application integration (EAI).** EAI technologies provide middleware that facilitates communication between multiple IT systems and business intelligence (BI) systems. EAI enables the BI system to subscribe to business events occurring within the enterprise.

- **Low-latency ETL micro batches.** It is difficult to extract data from source systems using only ETL; hence, low-latency ETL micro batches are often combined with CDC techniques to extract data from source systems in real time. They can also work with other extraction mechanisms (e.g., messages, services, replication).

- **Data federation.** Data federation fetches the data from multiple sources in real time. There is no data integration being done, but the disparate data sources (even geographically distributed ones) are transparently mapped and connected. Data federation can also use APIs to access the data. Sometimes this is called data virtualization.

Choosing the right technology depends upon the data requirements of the business scenarios the organization is trying to implement. Factors to consider include data volumes, frequency, data integrity, and transformation requirements.

Conclusion

Through this *Update*, we have articulated the need for an enterprise to have a near-real-time framework that will enable it to sense, analyze, and respond to opportunities and threats to its business. Gone are the days when an organization could afford to conduct these vital activities in an informal, intuitive, and individualistic style. An RTE framework needs to be formal, as objective as possible, and collaborative in order for an organization to be on top of its game.

About the Authors

Jagdish Bhandarkar is CTO of SLK Software. Previously, he was CTO of AspenWoods LLC and Principal/Head of the Enterprise Architecture Competency at Infosys, where he evangelized several programs in innovation, thought leadership, and technical skills for architects. He is the coauthor of the Big Data Spectrum *and* Next Generation IT Architecture *practitioner books. He holds an MS in software engineering.*

Lakshmanan G is Senior Principal at Infosys, where he is Head of the Community of Architects. Previously, he was Head of Infosys's EMEA IT Architecture Practice and its Engineering SaaS Practice. Lakshmanan is the coeditor of the Next Generation IT Architecture *and* Big Data Spectrum *practitioner books and has contributed to* Cutter Business Technology Journal*. Lakshmanan has a master of engineering degree.*

Raghavan S is Senior Director of BPO at Infosys. Previously, he was Product Manager for Infosys's Social Media platform, a SaaS offering, and Head of Infosys's technology consulting group for clients in Europe, the Middle East, and Africa. He has reviewed/defined the IT strategy of telecom service providers, banks, and retail organizations in the US, Europe, and Australia. Raghavan holds a bachelor's degree in electronics and communication engineering from the Bharathidasan University, Tamilnadu, India.

Shyam Doddavula is Associate VP and Principal Product Architect at Infosys, where he leads the blockchain CoE of the Infosys Center for Emerging Technologies. He has 20 years' experience in enterprise application development and technology R&D, looking into emerging tech trends, defining enterprise adoption strategies, and developing accelerators. He has over 10 invention disclosures and several publications.